

# An Introduction to ScaLAPACK



Tony Drummond  
Lawrence Berkeley National Laboratory  
[LADrummond@lbl.gov](mailto:LADrummond@lbl.gov)



# An Introduction to ScaLAPACK



## Outline:

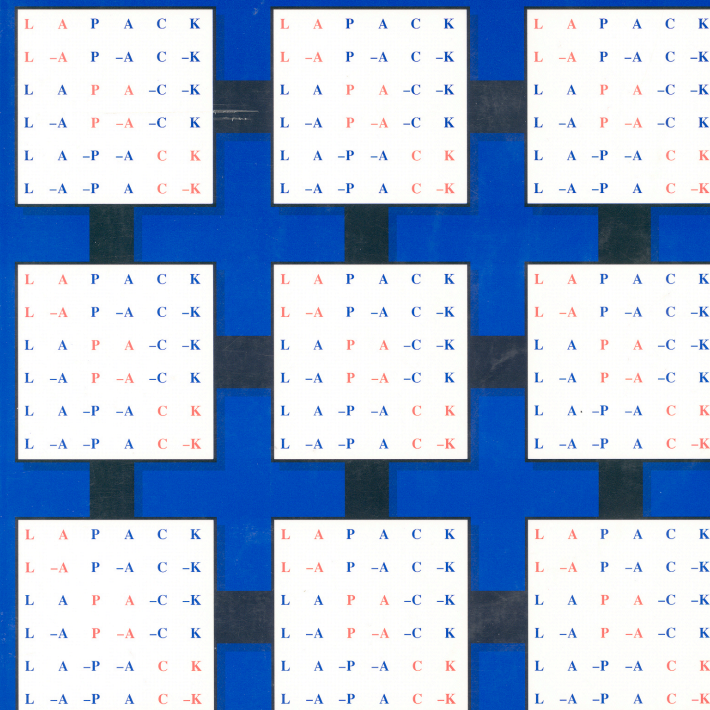
- Introduction
- Applications
- ScaLAPACK Functionality
- Software Hierarchy and Interfaces
- ScaLAPACK User Interface
- Performance
- Afternoon: Hands-ON

# An Introduction to ScaLAPACK



## ScaLAPACK Users' Guide

L. S. Blackford • J. Choi • A. Cleary • E. D'Azevedo  
J. Demmel • I. Dhillon • J. Dongarra • S. Hammarling  
G. Henry • A. Petitet • K. Stanley • D. Walker • R. C. Whaley



## Team of Developers:

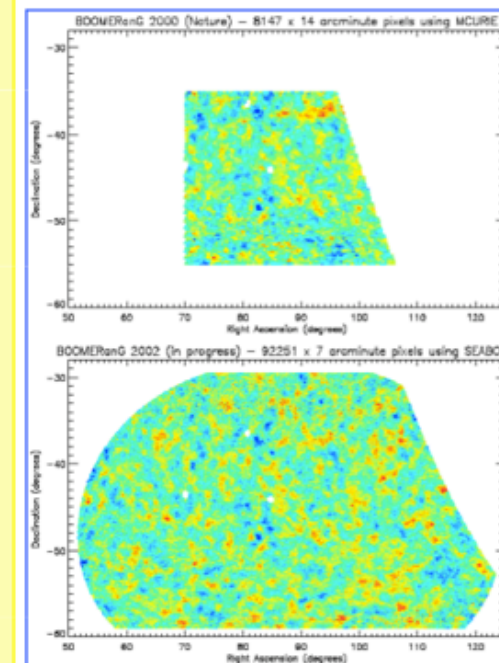
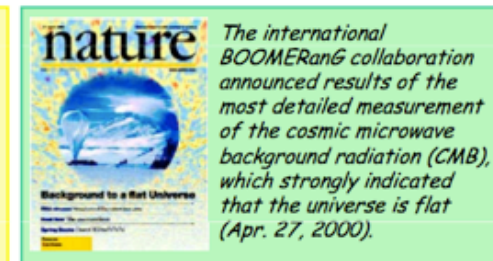
- Susan Blackford
- Jaeyoung Choi, Soongsil University
- Andy Cleary, LLNL
- Ed D'Azevedo, ORNL
- Jim Demmel, UCB
- Inderjit Dhillon, UT Austin
- Jack Dongarra, UTK
- Ray Fellers, LLNL
- Sven Hammarling, NAG
- Greg Henry, Intel
- Sherry Li, LBNL
- Osni Marques, LBNL
- Caroline Papadopoulos, UCSD
- Antoine Petitet, UTK
- Ken Stanley, UCB
- Françoise Tisseur, Manchester



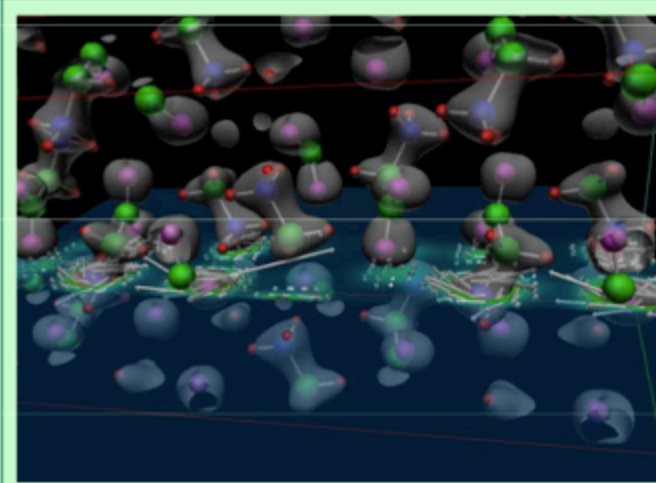
# MOTIVATION: Applications



- The statistics of the tiny variations in the CMB (the faint echo of the Big Bang) allows the determination of the fundamental parameters of cosmology to the percent level or better.
- MADCAP (Microwave Anisotropy Dataset Computational Analysis Package)
  - Makes maps from observations of the CMB and then calculates their angular power spectra. (See <http://crd.lbl.gov/~borrill>).
  - Calculations are dominated by the solution of linear systems of the form  $M=A^{-1}B$  for dense  $n \times n$  matrices  $A$  and  $B$  scaling as  $O(n^3)$  in flops. MADCAP uses **ScaLAPACK** for those calculations.
- On a Cray T3E (original code):
  - Cholesky factorization and triangular solve.
  - Typically reached 70-80% peak performance.
  - Solution of systems with  $n \sim 10^4$  using tens of processors.
  - The results demonstrated that the Universe is spatially flat, comprising 70% dark energy, 25% dark matter, and only 5% ordinary matter.
- On an IBM SP:
  - Porting was trivial but tests showed only 20-30% peak performance.
  - Code rewritten to use triangular matrix inversion and triangular matrix multiplication → one-day work !
  - Performance increased to 50-60% peak.
  - Solution of previously intractable systems with  $n \sim 10^5$  using hundreds of processors.

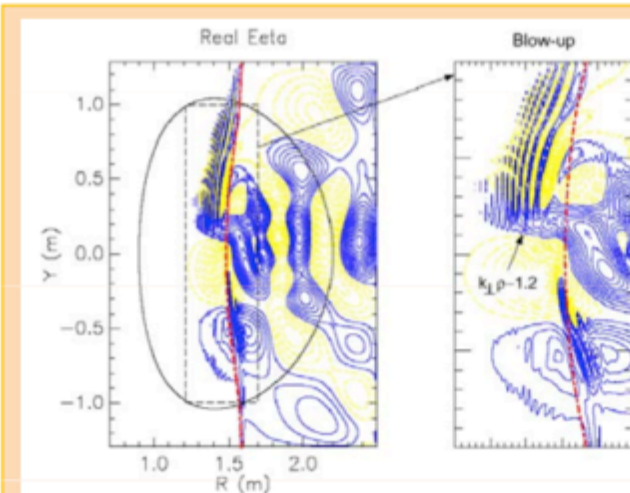
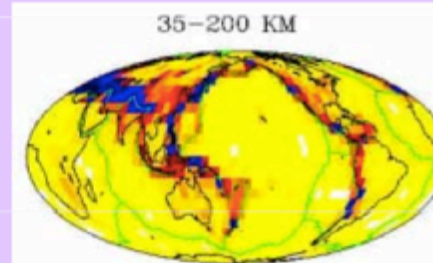


# MOTIVATION: Applications



Induced current (white arrows) and charge density (colored plane and gray surface) in crystallized glycine due to an external field; courtesy of Louie, Yoon, Pfrommer and Canning (UCB and LBNL).

Model for the internal structure of the Earth, resolution matrix (Vasco and Marques)



Two **ScaLAPACK** routines, **PZGETRF** and **PZGETRS**, are used for solution of linear systems in the spectral algorithms based **AORSA** code (Batchelor et al.), which is intended for the study of electromagnetic wave-plasma interactions. The code reaches 68% of peak performance on 1936 processors of an IBM SP.

Advanced Computational Research in Fusion (SciDAC Project, PI Mitch Pindzola). Point of contact: Dario Mitnik (Dept. of Physics, Rollins College). Mitnik attended the workshop on the ACTS Collection in September 2000 and afterwards actively used ACTS tools, in particular ScaLAPACK. Dario has worked on the development, testing and support of new scientific simulation codes related to the study of atomic dynamics using time-dependent close coupling lattice and time-independent methods. He has reported that this work could not be carried out in sequential machines and that ScaLAPACK was fundamental for the parallelization of these codes.

# An Introduction to ScaLAPACK



## Outline:

- Introduction
- Applications
- ScaLAPACK Functionality
- Software Hierarchy and Interfaces
- ScaLAPACK User Interface
- Performance
- Afternoon Hands-On

# ScaLAPACK's Functionality



$Ax = b$	Simple Driver	Expert Driver	Factor	Solve	Inversion	Conditioning Estimator	Iterative Refinement
Triangular				x	x	x	x
SPD	x	x	x	x	x	x	x
SPD Banded	x		x	x			
SPD Tridiagonal	x		x	x			
General	x	x	x	x	x	x	x
General Banded	x		x	x			
General Tridiagonal	x		x	x			
Least Squares	x		x	x			
GQR			x				
GRQ			x				
$Ax = \lambda x$ or $Ax = \lambda Bx$	Simple Driver	Expert Driver	Reduction	Solution			
Symmetric	x	x	x	x			
General	x	x	x	x			
Generalized BSPD	x		x	x			
SVD			x	x			

SPD: symmetric positive definite  
 GQR: generalized QR  
 GRQ: generalized RQ  
 BSPD: banded SPD  
 SVD: singular value decomposition

distinct computational task  
 standard types of problems



# An Introduction to ScaLAPACK

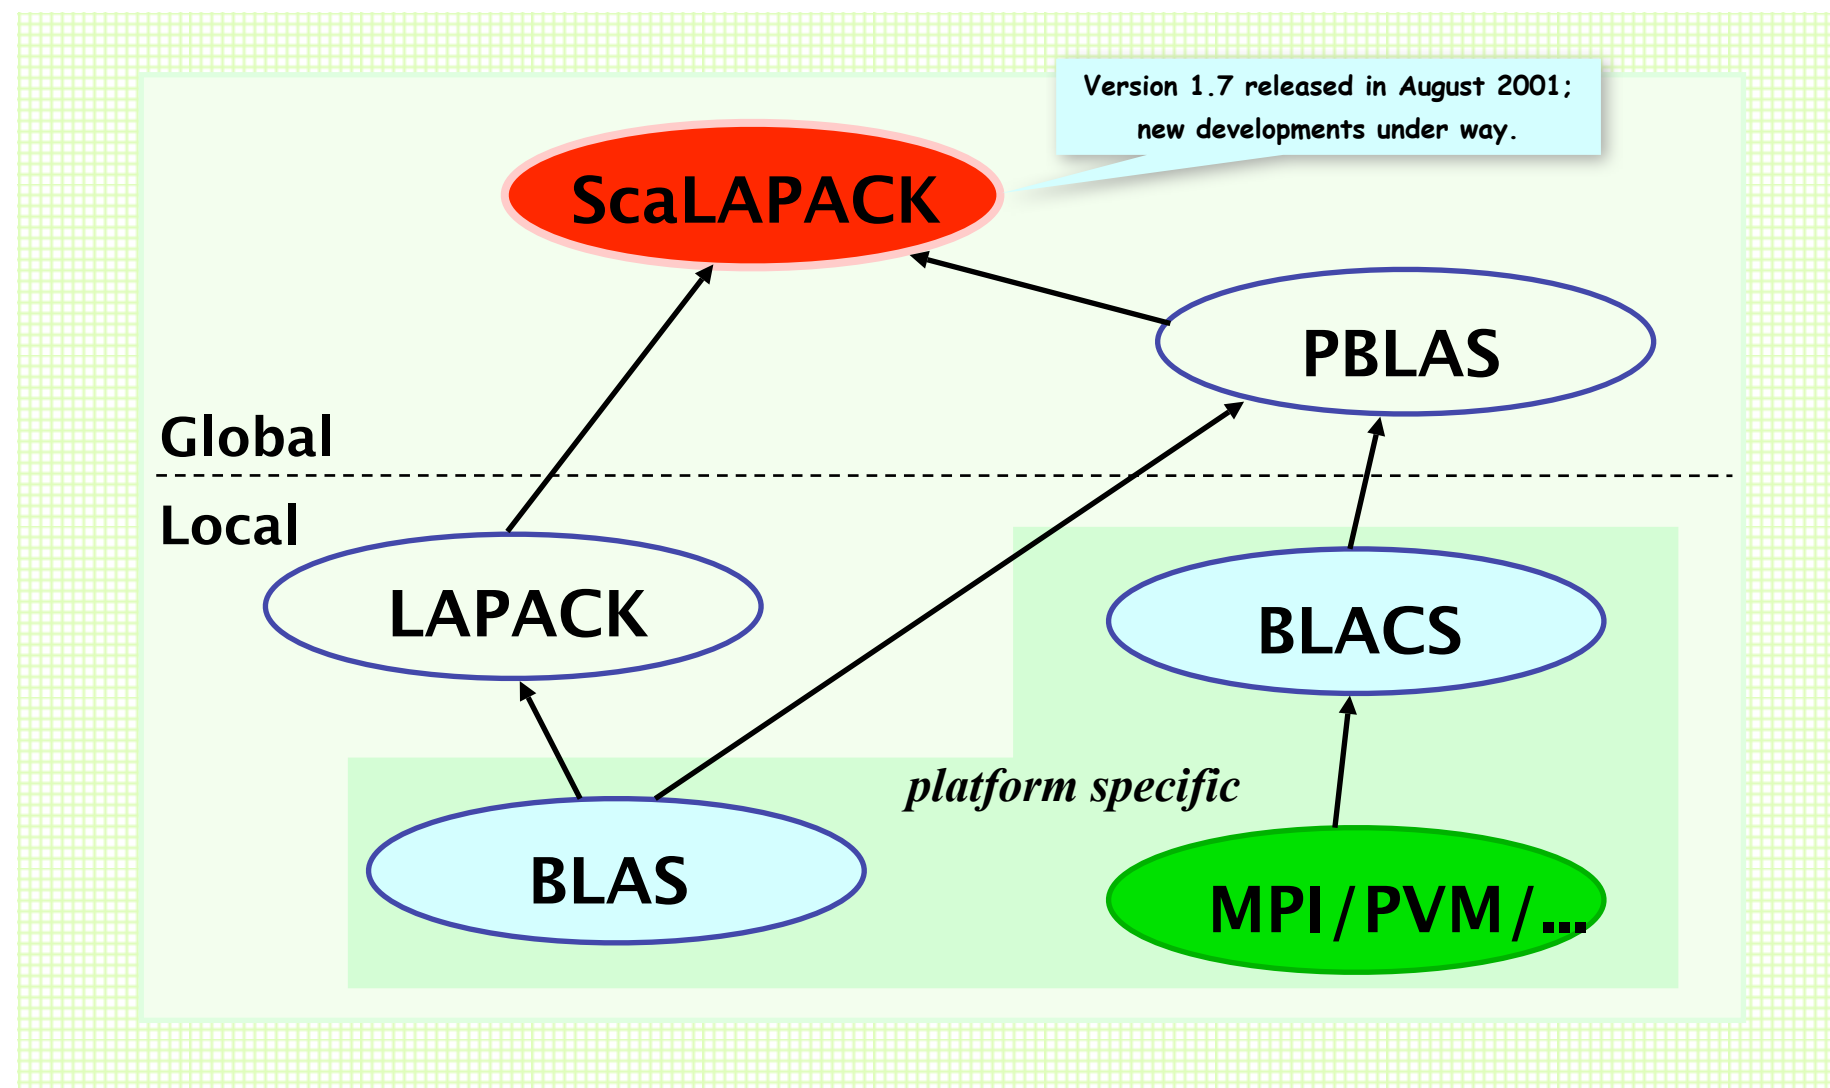


## Outline:

- Introduction
- Applications
- ScaLAPACK Functionality
- Software Hierarchy and Interfaces
- ScaLAPACK User Interface
- Performance
- Afternoon Hands-On



# ScaLAPACK's Software Structure

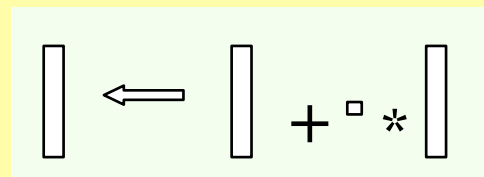


# BLAS: Basic Linear Algebra Subroutines

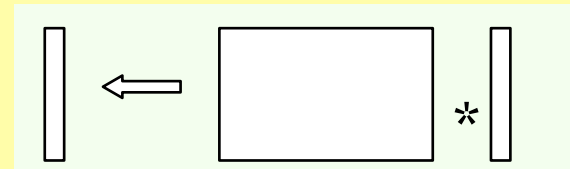


## BLAS LEVELS:

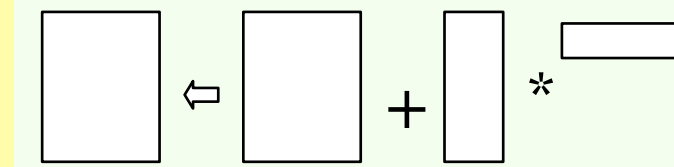
- Level 1 BLAS: vector-vector



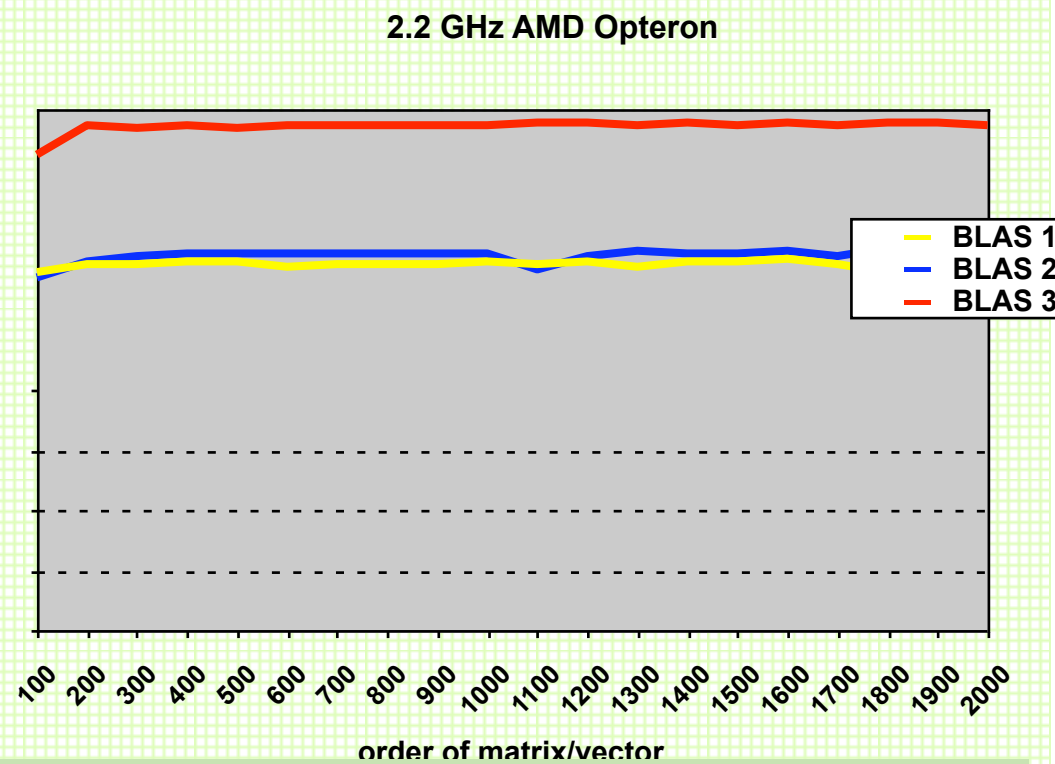
- Level 2 BLAS: matrix-vector



- Level 3 BLAS: matrix-matrix



Mflop/s



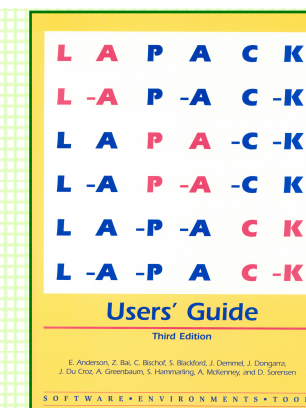
## Design Considerations:

- Portability
- Performance; development of blocked algorithms is important for performance!

# LAPACK: A Dense Linear Algebra Package



- Linear Algebra library written in Fortran 77 (Fortran 90)
- Combine algorithms from LINPACK and EISPACK into a single package.
- Efficient on a wide range of computers (RISC, Vector, SMPs).
- Built atop level 1, 2, and 3 BLAS Basic problems:
  - Linear systems:  $Ax = b$
  - Least squares:  $\min \|Ax - b\|_2$
  - Singular value decomposition:  $A = U\Sigma V^T$
  - Eigenvalues and eigenvectors:  $Az = \lambda z, Az = \lambda Bz$
- LAPACK does not provide routines for structured problems or general sparse matrices (i.e. sparse storage formats such as compressed-row, -column, -diagonal, skyline ...).



[netlib.org](http://netlib.org)

## BLACS: Basic Linear Algebra Communication Subroutines

- Response to Message Passing based distributed communications
- Associate widely recognized mnemonic names with communication operations. This improves:
  - program readability
  - self-documenting quality of the code.
- Promote efficiency by identifying frequently occurring operations of linear algebra which can be optimized on various computers.



# Basic Concepts of The BLACS Interface



- Promote efficiency by identifying common operations of linear algebra that can be optimized on various computers.
- Processes are embedded in a two-dimensional grid.

Example: a 3x4 grid

	0	1	2	3
0	0	1	2	3
1	4	5	6	7
2	8	9	10	11

- An operation which involves more than one sender and one receiver is called a scoped operation.

Scope	Meaning
Row	All processes in a process row participate.
Column	All processes in a process column participate.
All	All processes in the process grid participate.

# BLACS Communication Routines



Send/Receive:

```
_xxSD2D (ICTXT, [UPLO,DIAG],M,N,A,LDA,RDEST,CDEST)
_xxRV2D (ICTXT, [UPLO,DIAG],M,N,A,LDA,RSRC,CSRC)
```

<b>_ (Data type)</b>	<b>xx (Matrix type)</b>
I: Integer, S: Real, D: Double Precision, C: Complex, Z: Double Complex.	GE: General rectangular matrix TR: Trapezoidal matrix

Broadcast:

```
_xxBS2D (ICTXT, SCOPE, TOP, [UPLO,DIAG],M,N,A,LDA)
_xxBR2D (ICTXT, SCOPE, TOP, [UPLO,DIAG],M,N,A,LDA,RSRC,CSRC)
```

<b>SCOPE</b>	<b>TOP</b>
'Row' 'Column' 'All'	' ' (default) 'Increasing Ring' '1-tree' ...

## BLACS Context



- BLACS context  $\Leftrightarrow$  MPI communicator
- The BLACS context is the BLACS mechanism for partitioning communication space.
- A message in a context cannot be sent or received in another context.
- The context allows the user to
  - create arbitrary groups of processes
  - create multiple overlapping and/or disjoint grids
  - isolate each process grid so that grids do not interfere with each other

# An Example Code Using BLACS



```

      M
*  Get system information
  CALL BLACS_PINFO( IAM, NPROCS )
      M
*  Get default system context
  CALL BLACS_GET( 0, 0, ICTXT )
      M
*  Define 1 x (NPROCS/2+1) process grid
  NPROW = 1
  NPCOL = NPROCS / 2 + 1
  CALL BLACS_GRIDINIT( ICTXT, 'Row', NPROW, NPCOL )
  CALL BLACS_GRIDINFO( ICTXT, NPROW, NPCOL, MYROW, MYCOL )
*  If I'm not in the grid, go to end of program
  IF( MYROW.NE.-1 ) THEN
    IF( MYROW.EQ.0 .AND. MYCOL.EQ.0 ) THEN
      CALL DGESD2D( ICTXT, 5, 1, X, 5, 1, 0 )
    ELSE IF( MYROW.EQ.1 .AND. MYCOL.EQ.0 ) THEN
      CALL DGERV2D( ICTXT, 5, 1, Y, 5, 0, 0 )
    END IF
      M
    CALL BLACS_GRIDEXIT( ICTXT )
  END IF
      M
  CALL BLACS_EXIT( 0 )
END
```

*(out) uniquely identifies each process*  
*(out) number of processes available*

*(in) integer handle indicating the context*  
*(in) use (default) system context*  
*(out) BLACS context*

*(output)*  
*process row and*  
*column coordinate*

*send X to process (1,0)*

*receive X from process (0,0)*

*leave context*

*exit from the BLACS*



# PBLAS: Parallel BLAS



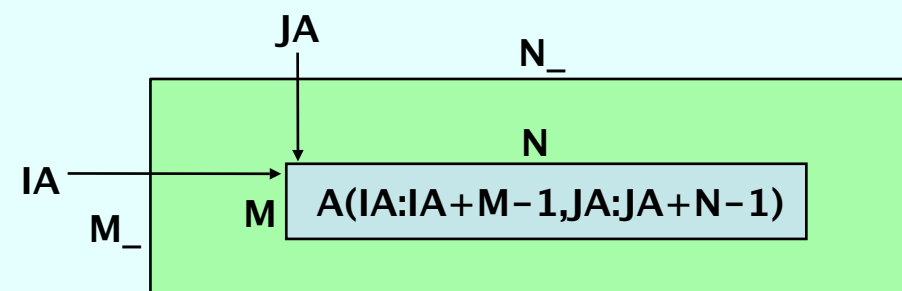
- Similar to the BLAS in portability, functionality and naming.
- Built atop the BLAS and BLACS
- Provide global view of matrix

```
CALL DGEXXX( M, N, A( IA, JA ), LDA, ... )
```

} BLAS

```
CALL PDGEXXX( M, N, A, IA, JA, DESCA, ... )
```

} PBLAS



Array descriptor (to be reviewed later)

# An Introduction to ScaLAPACK



## Outline:

- Introduction
- Applications
- ScaLAPACK Functionality
- Software Hierarchy and Interfaces
- ScaLAPACK User Interface
- Performance
- Afternoon Hands-ON

# ScaLAPACK Design Goals

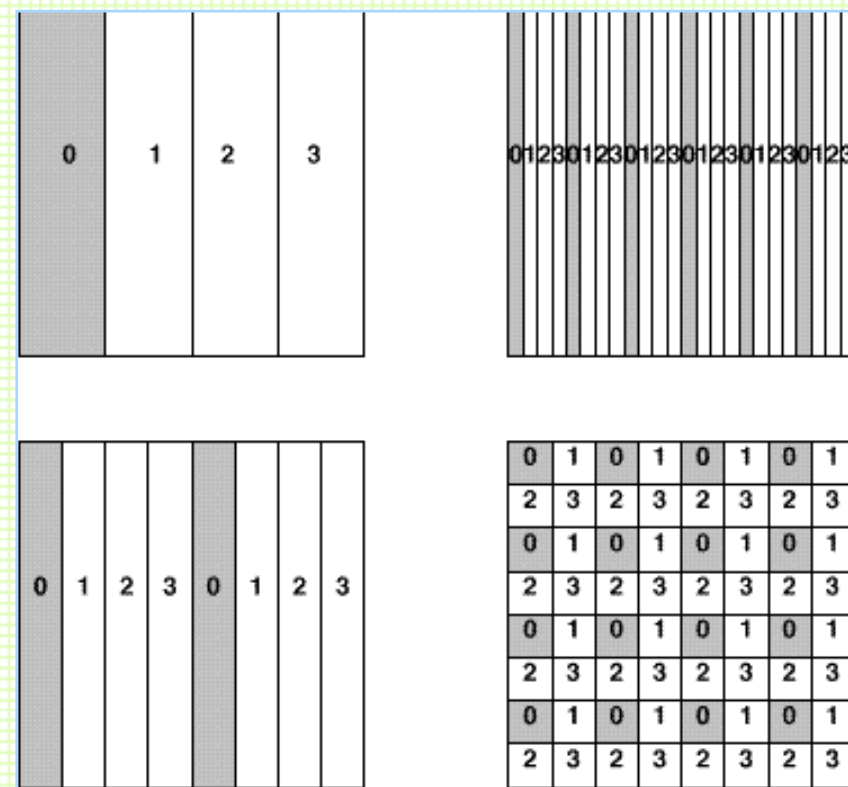


- Efficiency
  - Optimized computation and communication engines
  - Block-partitioned algorithms (Level 3 BLAS) for good node performance
- Reliability
  - Whenever possible, use LAPACK algorithms and error bounds.
- Scalability
  - As the problem size and number of processors grow
  - Replace LAPACK algorithm that did not scale (new ones into LAPACK)
- Portability
  - Isolate machine dependencies to BLAS and the BLACS
- Flexibility
  - Modularity: build rich set of linear algebra tools (BLAS, BLACS, PBLAS)
- Ease-of-Use
  - Calling interface similar to LAPACK

# ScaLAPACK: Data Layouts



- 1D block and column distributions
- 1D block-cycle column and 2D block-cyclic distribution
- 2D block-cyclic distribution used in ScaLAPACK for dense matrices

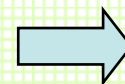




# How does 2D Block Cyclic Distribution Work



5x5 matrix partitioned in 2x2 blocks

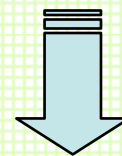
$$\begin{pmatrix} \mathbf{a_{11}} & \mathbf{a_{12}} & \mathbf{a_{13}} & \mathbf{a_{14}} & \mathbf{a_{15}} \\ \mathbf{a_{21}} & \mathbf{a_{22}} & \mathbf{a_{23}} & \mathbf{a_{24}} & \mathbf{a_{25}} \\ \mathbf{a_{31}} & \mathbf{a_{32}} & \mathbf{a_{33}} & \mathbf{a_{34}} & \mathbf{a_{35}} \\ \mathbf{a_{41}} & \mathbf{a_{42}} & \mathbf{a_{43}} & \mathbf{a_{44}} & \mathbf{a_{45}} \\ \mathbf{a_{51}} & \mathbf{a_{52}} & \mathbf{a_{53}} & \mathbf{a_{54}} & \mathbf{a_{55}} \end{pmatrix}$$


2x2 process grid point of view

$\mathbf{a_{11}}$ $\mathbf{a_{12}}$ $\mathbf{a_{15}}$ $\mathbf{a_{21}}$ $\mathbf{a_{22}}$ $\mathbf{a_{25}}$ $\mathbf{a_{51}}$ $\mathbf{a_{52}}$ $\mathbf{a_{55}}$	$\mathbf{a_{13}}$ $\mathbf{a_{14}}$ $\mathbf{a_{23}}$ $\mathbf{a_{24}}$ $\mathbf{a_{53}}$ $\mathbf{a_{54}}$
$\mathbf{a_{31}}$ $\mathbf{a_{32}}$ $\mathbf{a_{35}}$ $\mathbf{a_{41}}$ $\mathbf{a_{42}}$ $\mathbf{a_{45}}$	$\mathbf{a_{33}}$ $\mathbf{a_{34}}$ $\mathbf{a_{43}}$ $\mathbf{a_{44}}$

# An Example of 2D Block Cyclic Distribution

1.1	1.2	1.3	1.4	1.5
-2.1	2.2	2.3	2.4	2.5
-3.1	-3.2	3.3	3.4	3.5
-4.1	-4.2	-4.3	4.4	4.5
-5.1	-5.2	-5.3	-5.4	5.5



	0			1	
	a <sub>11</sub>	a <sub>12</sub>	a <sub>15</sub>	a <sub>13</sub>	a <sub>14</sub>
0	a <sub>21</sub>	a <sub>20</sub>	a <sub>25</sub>	a <sub>23</sub>	a <sub>24</sub>
	a <sub>51</sub>	a <sub>52</sub>	a <sub>55</sub>	a <sub>53</sub>	a <sub>54</sub>
1	a <sub>31</sub>	a <sub>32</sub>	a <sub>35</sub>	a <sub>33</sub>	a <sub>34</sub>
	a <sub>41</sub>	a <sub>42</sub>	a <sub>45</sub>	a <sub>43</sub>	a <sub>44</sub>

```

M
CALL BLACS_GRIDINFO( ICTXT, NPROW, NPCOL, MYROW, MYCOL )

IF ( MYROW.EQ.0 .AND. MYCOL.EQ.0 ) THEN
  A(1) = 1.1; A(2) = -2.1; A(3) = -5.1;
  A(1+LDA) = 1.2; A(2+LDA) = 2.2; A(3+LDA) = -5.2;
  A(1+2*LDA) = 1.5; A(2+3*LDA) = 2.5; A(3+4*LDA) =
ELSE IF ( MYROW.EQ.0 .AND. MYCOL.EQ.1 ) THEN
  A(1) = 1.3; A(2) = 2.3; A(3) = -5.3;
  A(1+LDA) = 1.4; A(2+LDA) = 2.4; A(3+LDA) = -5.4;
ELSE IF ( MYROW.EQ.1 .AND. MYCOL.EQ.0 ) THEN
  A(1) = -3.1; A(2) = -4.1;
  A(1+LDA) = -3.2; A(2+LDA) = -4.2;
  A(1+2*LDA) = 3.5; A(2+3*LDA) = 4.5;
ELSE IF ( MYROW.EQ.1 .AND. MYCOL.EQ.1 ) THEN
  A(1) = 3.3; A(2) = -4.3;
  A(1+LDA) = 3.4; A(2+LDA) = 4.4;
END IF

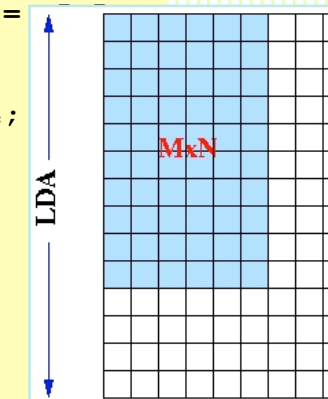
```

```

M
CALL PDGESVD( JOBU, JOBVT, M, N, A, IA, JA, DESCA, S, U, IU,
              JU, DESCU, VT, IVT, JVT, DESCVT, WORK, LWORK,
              INFO )
M

```

LDA is the leading  
dimension of the local  
array



## Why the headache of 2D block Cyclic Distribution?

- Ensures good load balance → performance and scalability (analysis of many algorithms to justify this layout).
- Encompasses a large number of data distribution schemes (but not all).
- Needs redistribution routines to go from one distribution to the other.
- See <http://acts.nersc.gov/scalapack/hands-on/datadist.html>

AID: <http://acts.nersc.gov/scalapack/hands-on/datadist.html>



ScaLAPACK Block Cyclic Data Distribution - Mozilla

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop

file:///C:/Documents%20and%20Settings/oam/ Search Print

### Block Cyclic Data Distribution

The block cyclic data distribution used by ScaLAPACK is justified by the analysis of many algorithms intended for linear algebra calculations and by the goal of achieving good load balancing. See [The Design of Linear Algebra Libraries for High Performance Computers](#), by J. Dongarra and D. Walker, and [Parallel Numerical Linear Algebra](#), by J. Demmel, M. Heath, and H. van der Vorst).

Use the the form below to specify the dimensions of a matrix, the desired blocking, and the process grid configuration. Then click on "distribute data": a new window will pop up with the corresponding block cyclic distribution. The acceptable values are indicated between parentheses; the defaults correspond to those of the matrix A used in [Exercise 4](#) in the [hands-on](#).

Note that although the PBLAS allow for non-square blocking factors, most ScaLAPACK routines do not, because of alignment constraints (which vary from routine to routine).

matrix dimensions	number of rows	<input type="text" value="9"/> (>0, <101)
	number of columns	<input type="text" value="9"/> (>0, <101)
matrix blocking	rows per block	<input type="text" value="2"/> (>0)
	columns per block	<input type="text" value="2"/> (>0)
process grid	rows	<input type="text" value="2"/> (>0)
	columns	<input type="text" value="3"/> (>0)

ScaLAPACK Tools Project Home Search

Block Cyclic Data Distribution

The table lists the data on each process. A = global array, B = local array, array indices start at 1.

Process (coordinates)	Array Values
0 (0,0)	B(1,1)=A(1,1) B(1,2)=A(1,2) B(2,1)=A(2,1) B(2,2)=A(2,2) B(1,3)=A(1,7) B(1,4)=A(1,8) B(2,3)=A(2,7) B(2,4)=A(2,8) B(3,1)=A(5,1) B(3,2)=A(5,2) B(4,1)=A(6,1) B(4,2)=A(6,2) B(3,3)=A(5,7) B(3,4)=A(5,8) B(4,3)=A(6,7) B(4,4)=A(6,8) B(5,1)=A(9,1) B(5,2)=A(9,2) B(5,3)=A(9,7) B(5,4)=A(9,8)
1 (0,1)	B(1,1)=A(1,3) B(1,2)=A(1,4) B(2,1)=A(2,3) B(2,2)=A(2,4) B(1,3)=A(1,9) B(2,3)=A(2,9) B(3,1)=A(5,3) B(3,2)=A(5,4) B(4,1)=A(6,3) B(4,2)=A(6,4) B(3,3)=A(5,9) B(4,3)=A(6,9) B(5,1)=A(9,3) B(5,2)=A(9,4) B(5,3)=A(9,9)
2 (0,2)	B(1,1)=A(1,5) B(1,2)=A(1,6) B(2,1)=A(2,5) B(2,2)=A(2,6) B(3,1)=A(5,5) B(3,2)=A(5,6) B(4,1)=A(6,5) B(4,2)=A(6,6) B(5,1)=A(9,5) B(5,2)=A(9,6)
3 (1,0)	B(1,1)=A(3,1) B(1,2)=A(3,2) B(2,1)=A(4,1) B(2,2)=A(4,2) B(1,3)=A(3,7) B(1,4)=A(3,8) B(2,3)=A(4,7) B(2,4)=A(4,8) B(3,1)=A(7,1) B(3,2)=A(7,2) B(4,1)=A(8,1) B(4,2)=A(8,2) B(3,3)=A(7,7) B(3,4)=A(7,8) B(4,3)=A(8,7) B(4,4)=A(8,8)
4 (1,1)	B(1,1)=A(3,3) B(1,2)=A(3,4) B(2,1)=A(4,3) B(2,2)=A(4,4) B(1,3)=A(3,9) B(2,3)=A(4,9) B(3,1)=A(7,3) B(3,2)=A(7,4) B(4,1)=A(8,3) B(4,2)=A(8,4) B(3,3)=A(7,9) B(4,3)=A(8,9)
5 (1,2)	B(1,1)=A(3,5) B(1,2)=A(3,6) B(2,1)=A(4,5) B(2,2)=A(4,6) B(3,1)=A(7,5) B(3,2)=A(7,6) B(4,1)=A(8,5) B(4,2)=A(8,6)

The color scheme shows the distribution of the global array on the computational grid. For the sake of resolution, colors are only displayed for less than 16 processes.

0	0	1	1	2	2	0	0	1
0	0	1	1	2	2	0	0	1
3	3	4	4	5	5	3	3	4
3	3	4	4	5	5	3	3	4
0	0	1	1	2	2	0	0	1
0	0	1	1	2	2	0	0	1
3	3	4	4	5	5	3	3	4
3	3	4	4	5	5	3	3	4
0	0	1	1	2	2	0	0	1



## ScaLAPACK: Array Descriptors



- Each global data object is assigned an array descriptor.
- The array descriptor:
  - Contains information required to establish mapping between a global array entry and its corresponding process and memory location (uses concept of BLACS context).
  - Is differentiated by the `DTYPE_` (first entry) in the descriptor.
  - Provides a flexible framework to easily specify additional data distributions or matrix types.
- User must distribute all global arrays prior to the invocation of a ScaLAPACK routine, for example:
  - Each process generates its own submatrix.
  - One processor reads the matrix from a file and send pieces to other processors (may require message-passing for this).

# Array Descriptor for Dense Matrices



DESC_()	Symbolic Name	Scope	Definition
1	DTYPE_A	(global)	Descriptor type DTYPE_A=1 for dense matrices.
2	CTXT_A	(global)	BLACS context handle.
3	M_A	(global)	Number of rows in global array A.
4	N_A	(global)	Number of columns in global array A.
5	MB_A	(global)	Blocking factor used to distribute the rows of array A.
6	NB_A	(global)	Blocking factor used to distribute the columns of array A.
7	RSRC_A	(global)	Process row over which the first row of the array A is distributed.
8	CSRC_A	(global)	Process column over which the first column of the array A is distributed.
9	LLD_A	(local)	Leading dimension of the local array.

# Array Descriptor for Narrow Band Matrices



DESC_( )	Symbolic Name	Scope	Definition
1	DTYPE_A	(global)	Descriptor type DTYPE_A=501 for 1 x Pc process grid for band and tridiagonal matrices block-column distributed.
2	CTXT_A	(global)	BLACS context handle.
3	N_A	(global)	Number of columns in global array A.
4	NB_A	(global)	Blocking factor used to distribute the columns of array A.
5	CSRC_A	(global)	Process column over which the first column of the array A is distributed.
6	LLD_A	(local)	Leading dimension of the local array. For the tridiagonal subroutines, this entry is ignored.
7	—	—	Unused, reserved.

## Array Descriptor for Right Hand Sides for Narrow Band Linear Solvers



DESC_( )	Symbolic Name	Scope	Definition
1	DTYPE_B	(global)	Descriptor type DTYPE_B=502 for $P_r \times 1$ process grid for block-row distributed matrices
2	CTXT_B	(global)	BLACS context handle
3	M_B	(global)	Number of rows in global array B
4	MB_B	(global)	Blocking factor used to distribute the rows of array B
5	RSRC_B	(global)	Process row over which the first row of the array B is distributed
6	LLD_B	(local)	Leading dimension of the local array. For the tridiagonal subroutines, this entry is ignored
7	—	—	Unused, reserved

## ScaLAPACK Routines



Three types of routines:

- Driver routines: each of which solves a complete problem, for example, solving a system of linear equations or computing the eigenvalues of a real symmetric matrix.
- Computational routines: each of which performs a distinct computational task, for example an LU factorization or the reduction of a real symmetric matrix to tridiagonal form. Each driver routine calls a sequence of computational routines. Global and local input error-checking are performed for these routines.
- Auxiliary routines: which in turn can be classified as follows:
  - routines that perform subtasks of block-partitioned algorithms -- in particular, routines that implement unblocked versions of the algorithms; and
  - routines that perform some commonly required low-level computations, for example, scaling a matrix, computing a matrix-norm, or generating an elementary Householder matrix.



# ScaLAPACK Functionality



$\mathbf{Ax} = \mathbf{b}$	Driver type		Factor	Solve	Inversion	Conditioning estimator	Iterative Refinement
	Simple	Expert					
Triangular	x			x	x	x	x
SPD	x	x	x	x	x	x	x
SPD Banded	x		x	x			
SPD Tridiagonal	x		x	x			
General	x	x	x	x			
General Banded	x		x	x	x	x	x
General Tridiagonal	x		x	x			
Least Squares			x				
GQR	x		x	x			
GRQ			x				
$\mathbf{Ax} = \lambda \mathbf{x}$ or $\mathbf{Ax} = \lambda \mathbf{Bx}$	Simple	Expert	Reduce	Solve			
Symmetric	x		x	x			
General	x	x	x	x			
Generalized BSPD	x	x	x	x			
SVD			x	x			

## ScaLAPACK: Error Handling



- Driver and computational routines perform global and local input error-checking.
  - Global checking → synchronization
  - Local checking → validity
- No input error-checking is performed on the auxiliary routines.
- If an error is detected in a PBLAS or BLACS routine program execution stops.

## ScaLAPACK: Debugging Hints



- Look at ScaLAPACK example programs.
- Always check the value of INFO on exit from a ScaLAPACK routine.
- Query for size of workspace, LWORK = -1.
- Link to the Debug Level 1 BLACS (specified by BLACSDBGVL=1 in Bmake.inc).
- Consult errata files on netlib:  
<http://www.netlib.org/scalapack/errata.scalapack>  
<http://www.netlib.org/blacs/errata.blacs>

# An Introduction to ScaLAPACK



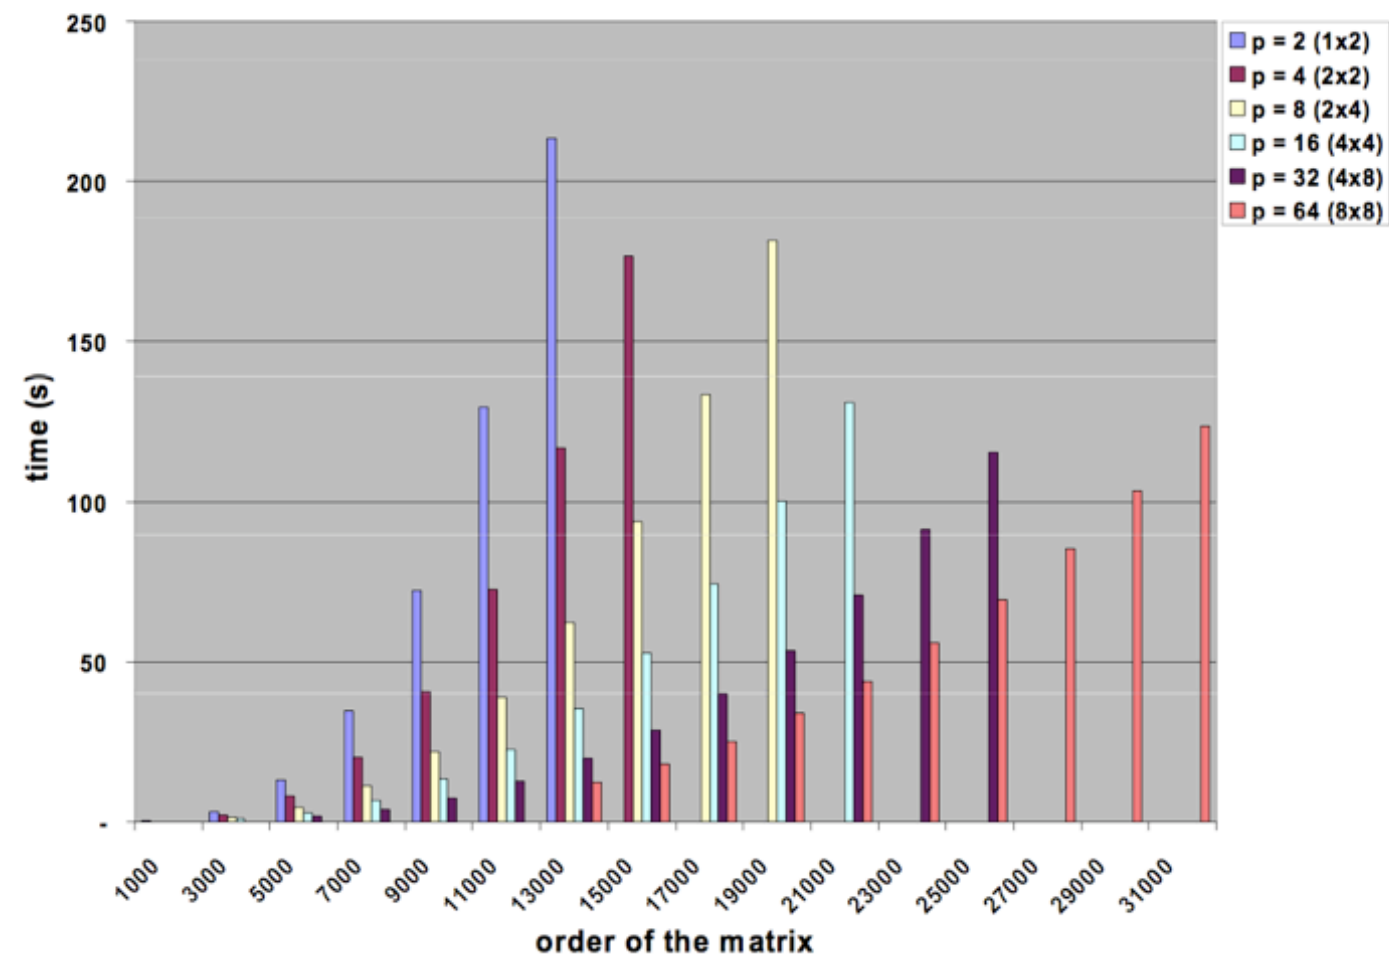
## Outline:

- Introduction
- Applications
- ScaLAPACK Functionality
- Software Hierarchy and Interfaces
- ScaLAPACK User Interface
- Performance
- Afternoon Hands-ON

# ScaLAPACK Performance



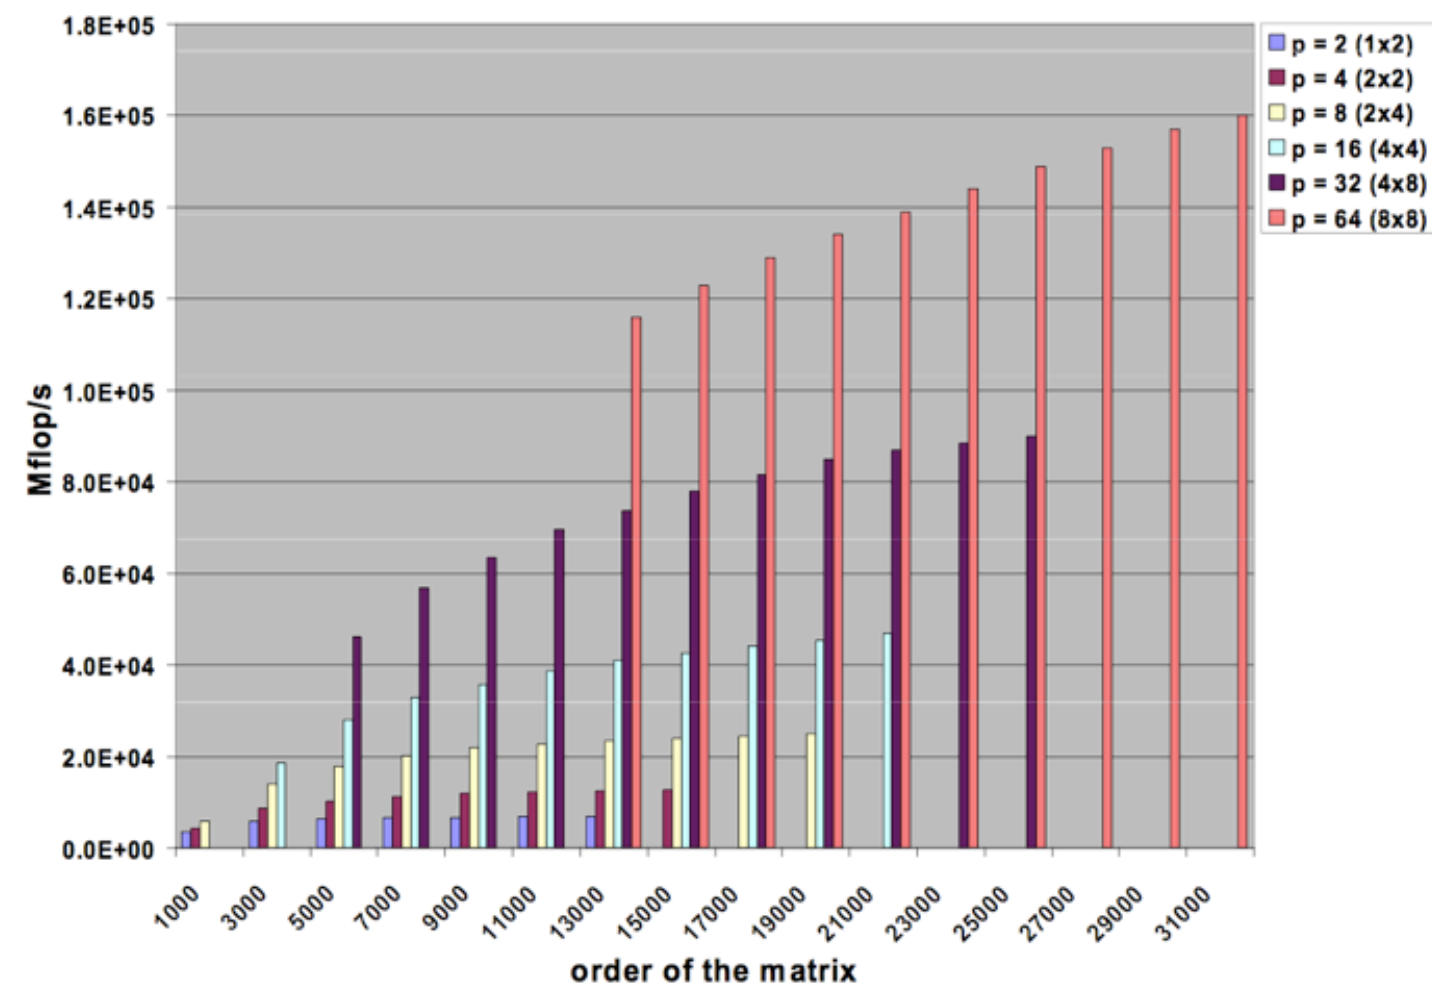
*LU on 2.2 GHz AMD Opteron (4.4 GFlop/s peak performance)*



# ScaLAPACK Performance



*LU+solve on 2.2 GHz AMD Opteron (4.4 GFlop/s peak performance)*

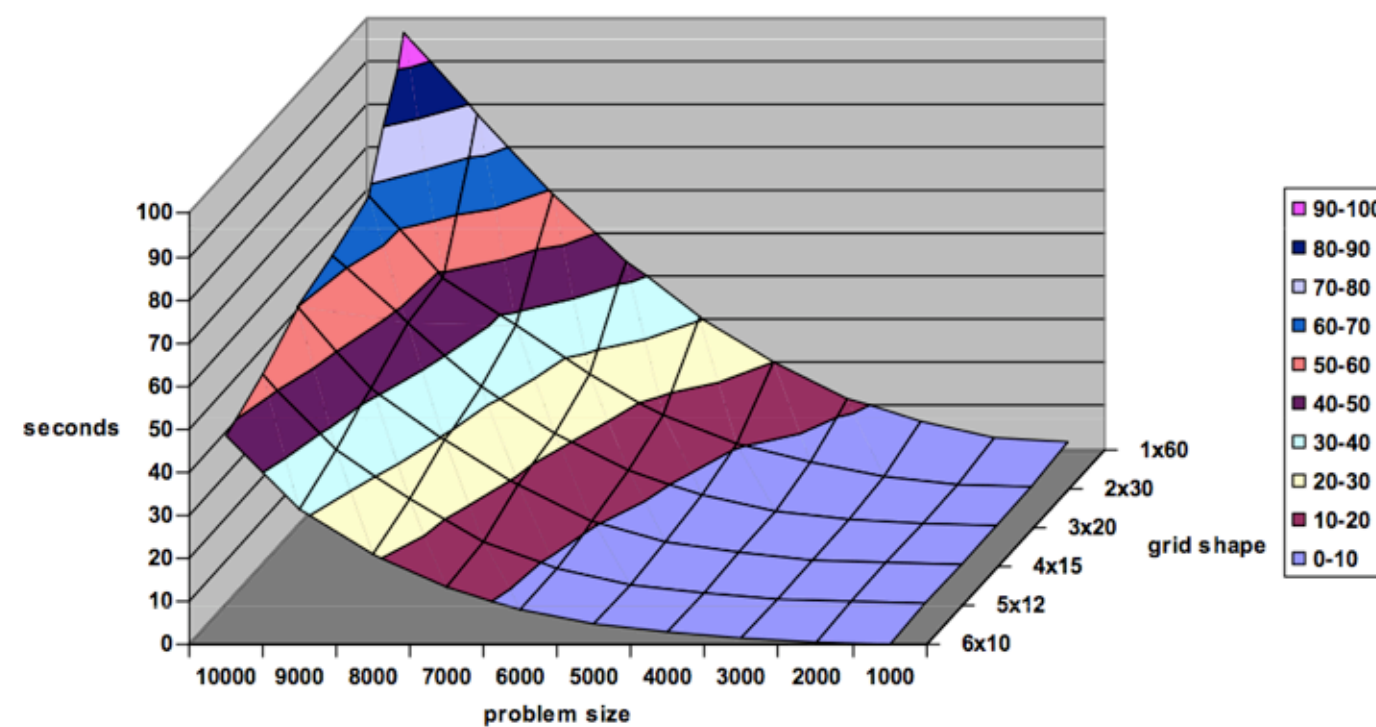




# ScaLAPACK Performance



Execution time of PDGESV for various grid shape



60 processors, Dual AMD Opteron 1.4GHz Cluster with Myrinet Interconnect, 2GB Memory

## Commercial use of ScaLAPACK



ScaLAPACK has been incorporated in the following commercial packages:

- Fujitsu
- Hewlett-Packard
- Hitachi
- IBM Parallel ESSL
- NAG Numerical Library
- Cray LIBSCI
- NEC Scientific Software Library
- Sun Scientific Software Library
- Visual Numerics (IMSL)

## Summary ScaLAPACK



- Library of high performance dense linear algebra routines for distributed-memory computing
- Reliable scalable and portable
- Reliable, scalable and portable
- Calling interface similar to LAPACK
- New developments on the way
- LAPACK/ScaLAPACK Forum:  
<http://icl.cs.utk.edu/lapack-forum>
- ScaLAPACK Users' Guide: <http://www.netlib.org/scalapack/slug>